BARGAINING WITH MIDDLEMEN

RANDALL A. BERRY, THÀNH NGUYEN, VIJAY SUBRAMANIAN
NORTHWESTERN UNIVERSITY

Abstract. In this paper, we consider a dynamic and decentralized market modeled by a non-cooperative networked bargaining game. Our goal is to study how the network structure of the market and the role of middlemen influence the market's efficiency and fairness. We introduce the concept of limit stationary equilibrium in a general trading network and use it to analyze how endogenous delay emerges in trade and how surplus is shared between sellers and buyers.

1. Introduction

In many markets, trade does not involve simply sellers and buyers but also one or more middlemen serving as intermediaries. For example, brokers and market makers fill this role in financial markets. Often, different buyers and sellers may not have access to the same set of middlemen due, for example, to various institutional or physical barriers. Such relationships are naturally modeled via a network. Understanding how the structure of this network effects the resulting market has been a topic attracting increased interest. This is partly due to the fact that on one hand, such questions are not adequately answered by classical economic models, while on the other hand phenomena, such as the recent financial crisis, suggest the importance of answering them.

A key question for network markets is explaining how prices form in such settings, which is a basic function of middlemen. Classic approaches such as competitive equilibrium analysis abstract away such questions as pointed out in the following quote from [15]:

> *Despite the important role played by intermediation in most markets, it is largely ignored by the standard theoretical literature. This is because a study of intermediation requires a basic model that describes explicitly the trade frictions that give rise to the function of intermediation. But this is missing from the standard market models, where the actual process of trading is left un-modeled.*

In recent years there has been a growing literature addressing such concerns including [16, 12, 13, 9, 14, 4]. This paper adds to this line of literature by studying non-cooperative bargaining in general network with middlemen similar to the network markets considered in [4]. Here, instead of price-setting agents, we consider non-cooperative bargaining: middlemen do not have full bargaining power as in [4]. In particular, we incorporate the non-cooperative bargaining model of [14] and add to it elements of search friction and generalize the trade network by allowing multiple trade routes. In this model, each node in a trade network consists of a population of agents. Non-cooperative bargaining occurs between pairs of agents that in adjacent nodes. As in [14], we study study the agents' behavior in the limit of large population sizes via the notion of a *limit stationary equilibrium.*[1]

We show existence of limit stationary equilibrium, and use this concept to investigate the efficiency of market, how bargaining with middlemen cause endogenous delay in equilibrium, and how network structure cause imbalance/unfairness in the share of surplus between sellers and buyers. These properties cannot be captured without non-cooperative and dynamic bargaining models.

[1]Further comparison with the existing literature is discussed in Section 3, where the solution concept of limit stationary equilibrium is introduced.

Our model and results give new economic insights into the combination of micro market mechanisms with a network. In particular, on how sunk-cost problems cause delays in trade for a dynamic model. This should be contrasted with static models in classical theory that predict no trade. In comparison to [14] where the absence of search friction can cause nonexistence of stationary market equilibria, here we show how slowing down trade with search friction can reestablish stationary equilibria. We show that with the presence of search friction, stationary equilibrium can be sustained, even with the possibility of multiple trade routes where the cheapest trade routes are preferred. This might suggest a direction for further research on the impact of trading speed on the dynamic behavior of an economy. However, inefficiency persists, in particular, in the new bargaining model, the strategy in which agents trade immediately whenever they meet a potential trading partner cannot be sustained at equilibrium unless the value of the good is high enough for the buyer; trade sometimes is delayed. Our results also reveal several interesting and non-monotonic properties of the equilibrium.

The remaining of the paper is organized as follows. Section 2 introduces the baseline noncooperative bargaining model, Section 3 discusses the solution concept of limit stationary equilibrium, Section 4 examines comparative analysis in some simple networks, and Section 5 concludes.

2. The Model

In this section we introduce the model that we will use.

Trading Network. We consider a group of sellers, buyers and middlemen interconnected by an underlying trading network, which is modeled as a directed graph, $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ (see Figure 1). Each node $i \in \mathcal{V}$ represents a population of N_i agents, all of which are either buyers, sellers or middlemen. Hence, we can partition the set of vertices into the following three disjoint sets: a set of sellers denoted by $\mathcal{I}$, a set of middlemen denoted by $\mathcal{J}$, and a set of buyers denoted by $\mathcal{K}$. An agent from the population at a node i will sometime be referred to as a type i agent. Trade "flows" over directed edges, i.e., a directed edge $(i, j) \in \mathcal{E}$ indicates that a type i agent can potentially directly trade with any type j agent. With a slight abuse of terminology, we often refer to two such agents as being connected by the edge (i, j). For a buyer to acquire a good from a seller, there must be a (directed) path from the buyer to the seller. If this path has length 1 then the two can directly trade, otherwise they must rely on middlemen to facilitate the trade. For simplicity, we consider networks in which any path between a buyer and seller contains at most one middleman, i.e all such paths are either length 1 or 2. An example of such a network is shown in Figure 1. With this assumption, the set of directed edges, $\mathcal{E}$ can also be partitioned into three disjoint sets: those that directly connect sellers to buyers denoted by $\mathcal{E}_1$, those that connect sellers to middlemen denoted by $\mathcal{E}_2$, and those that connect middlemen to buyers, denoted by $\mathcal{E}_3$.

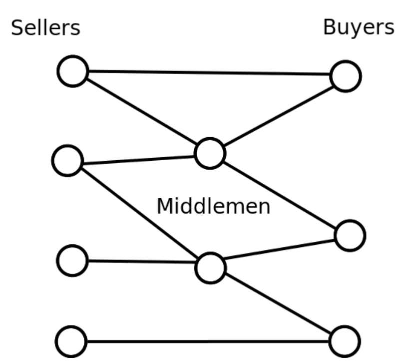

Figure 1. A network among sellers, buyers and middlemen

We assume that there is one type of indivisible good in this economy. All sellers produce identical goods and all buyers want to acquire these goods. The value that each buyer of type $k \in \mathcal{K}$ gets from an item is $V_k \geq 0$. At every period each agent can hold at most one unit of the good (an item). Thus, every time period, a middleman either has an item or does not have one. Hence, if there is a directed edge from node i to node j, a specific agent of type i can only trade with an agent of type j if the type i agent has a copy of the good and the type j agent does not; we refer to such a pair of agents as *feasible trading partners*. Note that sellers are assumed to always have a good available to trade and buyers are always willing to purchase a good. So for example, any two agents connected by an edge in the set E_1 are always feasible trading partners. For every edge $(i, j) \in \mathcal{E}$, we associate a non-negative transaction cost $C_{ij} \geq 0$; this cost is incurred when trade occurs between an agent at node i and one at node j.

The Bargaining Process. We consider an infinite horizon, discrete time repeated bargaining game, where agents discount their payoff at rate $0 < \delta < 1$. (The model can be extended to allow for heterogeneous discount rates.) Each period has multiple steps and is described as follows.

Step 1. One among all pairs of directly connected nodes $(i, j) \in \mathcal{E}$ is selected at random with a predetermined probability distribution $\pi(i, j)$ on the set of edges $\mathcal{E}$ and one node from each of the corresponding populations is selected uniformly at random. One of these agents is further selected to be a proposer (again chosen at random).

Step 2. If the agents are not feasible trading partners, then the game moves to the next period and restarts at step 1. Recall that this will occur if neither agent has the good or if both have the good.

Step 3. The proposer makes a take-it-or-leave it offer of a price at which he is willing to trade. If the trading partner refuses, the game moves to the next period. Otherwise, the two agents trade: one agent gives the item to and receives the money from the other, and the proposer pays for the transaction cost C_{ij}.[2] If a buyer or seller participates in a trade, they exit the game and are replaced by a clone. On the other hand, middlemen are long lived and do not produce nor consume; they earn money by flipping the good.

Step 4. The game moves to the next period, which starts from Step 1.

The game is denoted by $\Gamma(\mathcal{G}, \vec{C}, \vec{V}, \vec{N}, \delta)$, where $\vec{C}$ denotes the vector of links costs, $\vec{V}$ denotes the vector of buyer valuations and $\vec{N}$ denotes the vector of population sizes at each node i. Sometimes, we will simply refer to this game as Γ.

Remark. By appropriately choosing the distribution π and the choice of proposing agent when trade takes place, we can equivalently view the dynamics from the perspective of the agents such that the agents are picked independently (following some distribution) to be proposers and depending on the state of the agent (i.e., if the agent possesses the good or not), one among the appropriate edges is chosen following a distribution. Note that there is a possibility that the proposing agent might pick an edge along which no trade is possible owing to the picked agent having the same state as the proposing agent. This leads to search friction and can be contrasted with the model in [14] in which the proposing agent is always able to find a feasible trading partner if one exists. We prefer to model the dynamics from the perspective of edges as it is more general and fully subsumes the node perspective.

Replicated Economy. Given the bargaining game $\Gamma(\mathcal{G}, \vec{C}, \vec{V}, \vec{N}, \delta)$, the game's replications are defined as a game of the same structure, except the population size is increased by a factor of m, and the time gap between consecutive periods is reduced by a factor of T_m. Formally this is defined as follows:

[2] Actually, it does not matter who pays for this transaction cost, because the transaction cost is reflected in the proposing price.

Definition 2.1. *Given the game $\Gamma(\mathcal{G}, \vec{C}, \vec{V}, \vec{N}, \delta)$ and $m, T_m \in \mathbb{N}_+$, let $\delta' = \delta^{1/T_m}$. Then the (m, T_m)-replication of Γ, denoted by $\Gamma^m_{T_m}(\mathcal{G}, \vec{C}, \vec{V}, \vec{N}, \delta)$ is defined as $\Gamma(\mathcal{G}, \vec{C}, \vec{V}, m\vec{N}, \delta')$.*

Remark. The changing of discount rate δ is commonly used in the study of dynamical systems. It is clear that without changing δ, in the replicated economy each agent will need to wait for a longer and longer time to get selected, and thus his pay-off approaches 0. If initially each period takes one unit of time, then note that changing the discount rate to $\delta' = \delta^{1/T_m}$ is mathematically equivalent to changing the time gap between periods to become $1/T_m$ time units and keeping the discount rate fixed. Hence, for example, if we choose $T_m = c \cdot m$, it means we keep the rate that each agent see trading opportunities on the same order as in the original finite game. On the other hand $T_m >> m$ models a setting in which the rate at which agents trade at is increasing. In this paper, for simplicity, we will focus on the case $T_m = m$. Other choices of T_m do not affect our results, qualitatively.

3. Solution Concept and its Existence

Next we turn to the solution concept considered in this paper, which following [14] we call a *limit stationary equilibrium*. This is an equilibrium in which all players employ *stationary strategies*, defined as follows:

Definition 3.1. *A strategy profile (possibly mixed strategy) is called a stationary strategy if it only depends on an agent's identity, his state (owning or not owning and item) and the play of the game (which agent he is bargaining with, who is the proposer and what is proposed). More precisely, for example, assume that agent i and agent j are selected to bargain, and assume i owns an item, j does not, furthermore i is the proposer, then in this case a stationary strategy of the agent i is a distribution of proposing prices to agent j and a stationary strategy of the agent j is a probability of accepting the offer.*

Loosely, a limit stationary equilibrium in a profile of stationary strategies with two properties:

(1) Each agent's stationary strategy maximizes their expected pay-off assuming given probabilities μ_j for all $j \in \mathcal{J}$, which indicate the probability that a middleman selected from the population at node j owns a good in any period (equivalently, μ_j can be viewed as the steady-state fraction of middlemen owning a good).
(2) The assumed probabilities are required to be consistent with the given stationary strategies in the limiting replicated game as m increases without bound.

Such an equilibrium can be viewed as a type of fulfilled or rational expectations equilibrium (see e.g., [10]) in that agents can be viewed as making decisions based on a belief about the stationary probabilities (property 1) and these beliefs are required to be consistent with their resulting actions (property 2). Note that without taking the limit of large m, assuming such time-invariant probabilities is clearly not reasonable; for example, whenever a trade from i to j occurs it would increase the fraction of nodes at j holding an item and thus change this probability. However, in the limit of replicating the game such effects become negligible and as shown in the following, such equilibria do exist. A limit stationary equilibrium is also similar to the equilibrium concept used in mean-field games [1, 6, 7, 8, 11] in that the agents react to the distribution of actions (here possessing a good or not) of all the other agents. However, in contrast to the literature on mean-field games, here the agents are heterogenous with different agents reacting to the distribution of specific types of agents.

We will precisely define this equilibrium in this section. Before doing this, we will consider the first of the preceding properties and derive incentive constraints that an agent's stationary strategy needs to satisfy. We then turn to the second property, which involves considering a fluid limit of an underlying Markov process. After doing this, we combine these considerations to define this solution concept and further show that such an equilibrium always exists.

3.1. **Incentive Constraints.** To derive the needed incentive constraints, assume that each agent chooses an optimal stationary strategy to maximize their expected discounted pay-off given a set of probabilities $\{\mu_j : j \in \mathcal{J}\}$, which as discussed previously give the fraction of middlemen at each node j that hold an item at any time. In particular, in this setting the expected pay-off [3] of agent i will only depend on whether he has or does not have a good, which we denote by $u_0(i)$ and $u_1(i)$, respectively. Notice that because of the assumption that sellers and buyers exit the market after a successful trade, we have $u_0(i) = 0$ for every seller $i \in \mathcal{I}$; and $u_1(k) = V_k$ for every buyer $k \in \mathcal{K}$.

Consider the situation when edge (j, k) is chosen and a corresponding middleman of type $j \in \mathcal{J}$ who possesses an item is chosen to be the proposer to a buyer of type $k \in \mathcal{K}$. Again, abusing notation, we will refer to the specific chosen agents from the two populations as j and k. If the trade is successfully completed, then k possesses the item, thus agent j will demand from agent k the difference of the payoffs between the states before and after the trade (discounted by δ). Note that the state of j also changes, and therefore, if trade is successfully completed, then j's payoff is

$$\delta u_0(j) + \delta\big(u_1(k) - u_0(k)\big) - C_{jk}.$$

However, agent j has the option of not proposing a trade (or proposing something that will necessarily be rejected by the other party) and earn a payoff of $\delta u_1(j)$. For ease of exposition define the difference to be

$$z_{jk}(\delta) := \delta\Big(u_1(k) - u_0(k) - \big(u_1(j) - u_0(j)\big)\Big) - C_{jk}. \tag{1}$$

At equilibrium, the following properties, which we call the consistency conditions, will hold.

Definition 3.2 (Consistency conditions)**.**
(1) *If $\delta u_1(j) > \delta u_0(j) + \delta\big(u_1(k) - u_0(k)\big) - C_{jk}$, i.e., if $z_{jk}(\delta) < 0$, then agent j will never sell an item to agent k;*
(2) *If $\delta u_1(j) < \delta u_0(j) + \delta\big(u_1(k) - u_0(k)\big) - C_{jk}$, i.e., if $z_{jk}(\delta) > 0$, then agent j will sell the item to agent k with probability one whenever they're matched; and*
(3) *If $\delta u_1(j) = \delta u_0(j) + \delta\big(u_1(k) - u_0(k)\big) - C_{jk}$, i.e., if $z_{jk}(\delta) = 0$, then agent j is indifferent to the trade happening or not, so that the trade occurs with some probability $\lambda_{jk} \in [0, 1]$.*

From the third property above, it is clear that if trade between agents j and k occurs with probability $0 < \lambda_{jk} < 1$, then we must have $z_{jk} = 0$. Irrespective of whether the trade occurs or not, the payoff of agent j is

$$\delta u_1(j) + \max\{z_{jk}(\delta), 0\}.$$

Accounting for all events, the expected payoff of agent $j \in \mathcal{J}$, when possessing the good is

$$\sum_{k:(j,k)\in\mathcal{E}_3} \frac{\pi_{jk}}{2N_j}\Big(\delta u_1(j) + \max\{z_{jk}(\delta), 0\}\Big) + \left(1 - \sum_{k:(j,k)\in\mathcal{E}_3} \frac{\pi_{jk}}{2N_j}\right)\delta u_1(j).$$

Since agent $k \in \mathcal{K}$ departs as soon as she receives the item, and is replaced by a clone who does not have an item, with probability 1 agent j will find a feasible trading partner of type k. The expected payoff must satisfy the Bellman equation so that

$$u_1(j) = \sum_{k:(j,k)\in\mathcal{E}_3} \frac{\pi_{jk}}{2N_j}\Big(\delta u_1(j) + \max\{z_{jk}(\delta), 0\}\Big) + \left(1 - \sum_{k:(j,k)\in\mathcal{E}_3} \frac{\pi_{jk}}{2N_j}\right)\delta u_1(j). \tag{2}$$

After some algebraic manipulations, this is equivalent to

$$u_1(j) = \sum_{k:(j,k)\in\mathcal{E}_3} \frac{\pi_{jk}}{2N_j(1-\delta)} \max\{z_{jk}(\delta), 0\}. \tag{3}$$

[3] To avoid cumbersome notation we will use the same symbols for the payoffs for both finite m and $m \to \infty$.

Now consider the case of middleman $j \in \mathcal{J}$ that does not have an item and is the proposer. She has to trade with a seller $i \in \mathcal{I}$ to whom she is connected to. If the trade is successful, then payoff of agent j is

$$\delta u_1(j) - \delta\big(u_1(i) - u_0(i)\big) - C_{ij},$$

where $\delta\big(u_1(i) - u_0(i)\big)$ is the sum demanded by agent i. This has to be compared with $\delta u_0(j)$, the payoff for not trading at the current opportunity. Again define

$$z_{ji}(\delta) = \delta\Big(u_1(j) - u_0(j) - \big(u_1(i) - u_0(i)\big)\Big) - C_{ij} \tag{4}$$

so that the payoff of agent i is

$$\delta u_0(i) + \max\{z_{ji}(\delta), 0\}$$

where the value of $z_{ji}(\delta)$ determines whether trade happens.

Following the same steps as when agent j had the good, we get the expected payoff to be

$$u_0(j) = \sum_{i:(i,j)\in\mathcal{E}_2} \frac{\pi_{ij}}{2N_j(1-\delta)} \max\{z_{ji}(\delta), 0\}. \tag{5}$$

Using similar arguments, we can write the payoffs for every agent in our system as follows:

(1) Payoffs for sellers without an item are always 0, i.e., $u_0(i) \equiv 0$ for all $i \in \mathcal{I}$. Additionally, buyers with an item have their valuation as their payoff, i.e., $u_1(k) = V_k$ for all $k \in \mathcal{K}$;

(2) Seller $i \in \mathcal{I}$ is chosen and has an item to sell. Depending on the graph $\mathcal{G}$, the seller can trade with a middleman or directly with a buyer. The expected payoff is given by

$$u_1(i) = \sum_{k:(i,k)\in\mathcal{E}_1} \frac{\pi_{ik}}{2N_i(1-\delta)} \max\{z_{ik}(\delta), 0\} + \sum_{j:(i,j)\in\mathcal{E}_2} \frac{\pi_{ij}}{2N_i(1-\delta)} (1-\mu_j) \max\{z_{ij}(\delta), 0\} \tag{6}$$

where

$$z_{ik}(\delta) = \delta\Big(u_1(k) - u_0(k) - \big(u_1(i) - u_0(i)\big)\Big) - C_{ik}, \tag{7}$$

$$z_{ij}(\delta) = \delta\Big(u_1(j) - u_0(j) - \big(u_1(i) - u_0(i)\big)\Big) - C_{ij}. \tag{8}$$

Here, recall that N_i is the size of population at node i, and thus, for every $j \in \mathcal{J}$, $\frac{\pi_{ij}}{2N_i}(1-\mu_j)$ is the probability that conditional on holding an item, i is matched with j that does not hold a good and i is the proposer. Note that *search friction* impacts the transaction between the seller and the middleman;

(3) Buyer $k \in \mathcal{K}$ is chosen and does not have an item. Again, depending on the graph $\mathcal{G}$, the buyer can trade with a middleman who has a good or directly with a seller. We set $z_{ki}(\delta) = z_{ik}(\delta)$ and $z_{kj}(\delta) = z_{jk}(\delta)$. The expected payoff is

$$u_0(k) = \sum_{i:(i,k)\in\mathcal{E}_1} \frac{\pi_{ik}}{2N_k(1-\delta)} \max\{z_{ki}(\delta), 0\} + \sum_{j:(j,k)\in\mathcal{E}_3} \frac{\pi_{jk}}{2N_k(1-\delta)} \mu_j \max\{z_{kj}(\delta), 0\}. \tag{9}$$

Here too we incorporate *search friction*, so that given a buyer $k \in \mathcal{K}$, for every $j \in \mathcal{J}$, $\frac{\pi_{jk}}{2N_k}\mu_j$ is the probability that conditional on not holding an item, k is matched with j that holds a good and k is the proposer.

As the economy gets large, we need to consider the behavior of equations (3), (6), (5) and (9), where N_j is replaced by mN_j and δ is replaced by $\delta^{1/m}$, as m increases without bound. Since the payoffs are non-negative and are bounded by $\max_{k\in\mathcal{K}} V_k$, along subsequences, limits exist; by relabeling, if necessary, consider any such subsequence. We will now discuss properties of any such

a subsequence limit. First note that $\lim_{m\to\infty} m(1-\delta^{1/m}) = \ln(1/\delta)$. Then we get the following equations holding at each limit point:

$$\tag{10} \begin{aligned} \forall i \in \mathcal{I} \quad u_1(i) = & \sum_{k:(i,k)\in\mathcal{E}_1} \frac{\pi_{ik}}{2N_i \ln(1/\delta)} \max\{z_{ik}, 0\} \\ & + \sum_{j:(i,j)\in\mathcal{E}_2} \frac{\pi_{ij}}{2N_i \ln(1/\delta)} (1-\mu_j) \max\{z_{ij}, 0\}, \end{aligned}$$

$$\tag{11} \forall j \in \mathcal{J} \quad u_0(j) = \sum_{i:(i,j)\in\mathcal{E}_2} \frac{\pi_{ij}}{2N_j \ln(1/\delta)} \max\{z_{ji}, 0\},$$

$$\tag{12} \forall j \in \mathcal{J} \quad u_1(j) = \sum_{k:(j,k)\in\mathcal{E}_3} \frac{\pi_{jk}}{2N_j \ln(1/\delta)} \max\{z_{jk}, 0\},$$

$$\tag{13} \begin{aligned} \forall k \in \mathcal{K} \quad u_0(k) = & \sum_{i:(i,k)\in\mathcal{E}_1} \frac{\pi_{ik}}{2N_k \ln(1/\delta)} \max\{z_{ki}, 0\} \\ & + \sum_{j:(j,k)\in\mathcal{E}_3} \frac{\pi_{jk}}{2N_k \ln(1/\delta)} \mu_j \max\{z_{kj}, 0\}, \end{aligned}$$

where we have

$$\tag{14} z_{ik} = z_{ki} = \Big(u_1(k) - u_0(k) - \big(u_1(i) - u_0(i)\big)\Big) - C_{ik}$$

$$\tag{15} z_{ij} = z_{ji} = \Big(u_1(j) - u_0(j) - \big(u_1(i) - u_0(i)\big)\Big) - C_{ij};$$

$$\tag{16} z_{jk} = z_{kj} = \Big(u_1(k) - u_0(k) - \big(u_1(j) - u_0(j)\big)\Big) - C_{jk}.$$

Notice that by definition, as we consider the replicated game with larger and larger population, we also need to change the discount rate δ by $\delta^{1/m}$. As m approached infinity, $\delta^{1/m}$ becomes 1. As a result, (14)-(16) are obtained as limits of (1), (4) and (7)-(8).

We summarize the derivation by the following result.

Theorem 3.1. *Given $\{\mu_j : j \in \mathcal{J}\}$, every limit point of the replicated game Γ satisfies* (10)-(16).

3.2. **Limiting Trade Dynamics.** In the previous section we derived a set of constraints that an agent's stationary strategy would need to satisfy given that the fraction of middlemen owning an item at each node j is equal to a constant probability μ_j. Recall that such a strategy is characterized by a set of probabilities λ_{ij} for every edge $(i,j) \in \mathcal{E}$, which denotes the probability of trade occurring among a pair of feasible trading partners of type i and j. If $\lambda_{ij} = 1$, trade always occurs and if $\lambda_{ij} = 0$ it never occurs. Given any such a stationary strategy, the resulting dynamics can be viewed as a Markov process. In this section we consider the behavior of such a process for the replication of a game as m increases without bound. The main result is that for such processes, in the limit the fraction of middlemen at each node j does indeed converge to a fixed probability μ_j.

To begin suppose that there is a fixed fraction of middlemen μ_j at each node $j \in \mathcal{J}$ that hold the good for all time. We extend this to the sellers by setting $\mu_i \equiv 1$ for all $i \in \mathcal{I}$ and the buyers by setting $\mu_k \equiv 0$ for all $k \in \mathcal{K}$. The given values $\{\mu_j : j \in \mathcal{J}\}$ must obey a balance condition: for every node $j \in \mathcal{J}$, in every period the probability that the amount of goods held at node j increases by one or decreases by one should be equal. Mathematically, the balance condition can be written as

$$\tag{17} \sum_{k\in\mathcal{K}} \pi_{jk}\mu_j (1-\mu_k) \lambda_{jk} = \sum_{i\in\mathcal{I}} \pi_{ij}\mu_i (1-\mu_j) \lambda_{ij} \quad \forall j \in \mathcal{J}.$$

Here, for example, $\pi_{jk}\mu_j (1-\mu_k) \lambda_{jk}$ is the probability that trade occurs from j to k, which requires that link (j,k) is selected (with probability π_{jk}), that j and k are feasible trading partners (with

probability $\mu_j(1-\mu_k)$) and that trade occurs (with probability λ_{ij}). While the use of $\{\mu_i : i \in \mathcal{I}\}$ and $\{\mu_k : k \in \mathcal{K}\}$ above is for mathematical convenience, in more general networks where we allow middlemen to trade with each other, expressions similar to (17) will hold as the balance condition for every middlemen type where the terms will involve the state of other middlemen as well.

We will prove the existence of μ_j satisfying this balance condition by analyzing the Markov process that drives the state of the system, for a given set of stationary strategies $\{\lambda_{ij}\}$. Since the state of middlemen can change with time, the entire system can be represented by a vector-valued random process $\{X_j^m(t) : j \in \mathcal{J}\}_{t=1}^{\infty}$ where for the m replicated system we keep track of the number of agents who have the item at each middleman type $j \in \mathcal{J}$. For mathematical convenience, we will append $\{X_i^m(t) : i \in \mathcal{I}\}$ where $X_i^m(t) \equiv mN_i$ and $\{X_k^m(t) : k \in \mathcal{K}\}$ where $X_k^m(t) \equiv 0$ for the states of the sellers and the buyers, respectively. Since sellers exit the game as soon as they sell their good and are replaced by a clone with a good, at any given time any seller always possesses a good. A similar reasoning holds for the buyers never having a good.

For the m^{th} replication, the state transitions are given as follows for each $j \in \mathcal{J}$

$$X_j(t+1) = \begin{cases} \min(mN_j, X_j(t)+1) & \text{w. p. } \rho_j(+1) \\ \max(0, X_j(t)-1) & \text{w. p. } \rho_j(-1) \\ X_j(t) & \text{w. p. } 1-\rho_j(+1)-\rho_j(-1), \end{cases}$$

where

$$\rho_j(+1) = \left(1 - \frac{X_j(t)}{mN_j}\right) \sum_{i \in \mathcal{I}:(i,j)\in\mathcal{E}_2} \pi_{ij}\lambda_{ij}$$

is the probability that an agent of type j acquires a good in a given period, and

$$\rho_j(-1) = \frac{X_j(t)}{mN_j} \sum_{k \in \mathcal{K}:(j,k)\in\mathcal{E}_3} \pi_{jk}\lambda_{jk}$$

is the probability that a type j agent sells a good in a given period. As noted above, the states corresponding to sellers in $\mathcal{I}$ and buyers in $\mathcal{K}$ are fixed for all time.

This shows that we have a Markov process. Since the transition matrix satisfies Lipschitz conditions, we can analyze the fluid limit that is obtained by scaling time and space, i.e., by considering the process $\{\tilde{X}_v^m(t) : v \in \mathcal{V}\}$, where

$$\tilde{X}_v^m(t) := \frac{X_v^m(\lceil mt \rceil)}{m}, \qquad \forall v \in \mathcal{V}.$$

We will analyze the behavior of the process $\{\tilde{X}_v^m(t) : v \in \mathcal{V}\}_{t\in\mathbb{R}_+}$ when m increases without bound. Note that this the exact scaling considered by the replicated systems. We then have the following result.

Theorem 3.2. *Given a set of probabilities for trade $\{\lambda_{ij}, (i,j) \in \mathcal{E}\}$, the trading dynamic process described above converges to an unique state, which is the unique solution of* (17) *given by*

$$\forall i \in \mathcal{I}, \quad \mu_i = 1; \qquad \forall k \in \mathcal{K}, \quad \mu_k = 0; \tag{18}$$

$$\forall j \in \mathcal{J}, \quad \mu_j = \frac{\sum_{i\in\mathcal{I}:(i,j)\in\mathcal{E}_2} \pi_{ij}\lambda_{ij}}{\sum_{i\in\mathcal{I}:(i,j)\in\mathcal{E}_2} \pi_{ij}\lambda_{ij} + \sum_{k\in\mathcal{K}:(j,k)\in\mathcal{E}_3} \pi_{jk}\lambda_{jk}}. \tag{19}$$

Proof. See Appendix A.1. □

Note that $\{\mu_j : j \in \mathcal{J}\}$ is a continuous function of $\{\lambda_{ik} : (i,k) \in \mathcal{E}\}$. When middlemen are allowed to trade with each other, a subtler proof based on a fixed-point theorem can be used to show a similar result.

3.3. **Equilibrium.** Given the discussion about the incentive constraints and the convergence of the game dynamics above, we are now ready to define the solution concept of limit stationary equilibrium in this bargaining game. Loosely, as discussed previously, a limit stationary equilibrium is a profile of stationary strategies with two properties:

(1) Each agent's stationary strategy maximizes their expected pay-off assuming given probabilities μ_j for all $j \in \mathcal{J}$, which indicate the probability that a middleman selected from the population at node j owns a good in any period.
(2) The assumed probabilities are required to be consistent with the given stationary strategies in the limiting replicated game as m increases without bound. In particular, in the limiting game, for each agent, there will be two numerical values indicating the expected payoffs of agents at different state (owning/ not owning an item). A consistency requirement poses constraints between these payoffs and the stationary strategies similar to Definition 3.2.

Definition 3.3 (Limit Stationary Equilibrium)**.** *Given a stationary strategy, let* $0 \leq \lambda_{ij} \leq 1$ *be the overall probability that trade between* i *having a good and* j *wanting to buy one occur conditioned on the event that they are selected by the matching process. Let* μ *be the unique converging steady state of the random process defined by* λ_{ij}*, as discussed at the beginning of this section. Furthermore, let* $u_0(i), u_1(i)$ *be the expected payoff of agent* i *under the random process defined with* λ_{ij}*. This stationary strategy is a limit stationary equilibrium if*

(1) *Dynamic-state consistency:* μ *is the converging state of the dynamic defined with* λ*, that is* λ, μ *satisfy* (18)-(19)*;*
(2) *Payoff-state consistency:* $\lambda, \mu, \vec{u}$ *satisfy the incentive constraints defined in* (10)-(16)*; and*
(3) *Payoff-dynamic consistency: if* $z_{ij} > 0$ *then* $\lambda_{ij} = 1$*; if* $z_{ij} < 0$ *then* $\lambda_{ij} = 0$*; and if* $0 < \lambda_{ij} < 1$*, then* $z_{ij} = 0$ *for all links* (ij) *in the network* $\mathcal{G}$*, and are defined in* (14)-(16)*.*

We next show that a limit stationary equilibrium always exists. The proof of this theorem is based on the standard fixed-point theorem argument. For completeness, we provide such a detailed proof in the appendix. Notice that the equilibrium might not be unique; however, in the next section, we show several simple networks in which such an equilibrium is unique and exhibits interesting properties.

Theorem 3.3. *For a bargaining game,* $\Gamma(\mathcal{G}, \vec{C}, \vec{V}, \vec{N}, \delta)$*, a limit stationary equilibrium always exists.*

Proof. See Appendix A.2. □

4. Comparative Studies

In this section we will focus our comparative studies on when the discount factor δ approaches 1, in which case we sometimes refer to as when "agents are being patient" or "vanishing bargaining friction" [4]. We start by showing that despite the local bargaining set-up, a global efficiency emerges in the equilibria. In the rest of the section we give some comparative statics on how the network structure influences the payoffs and trade patterns in our model.

4.1. **Efficiency.**

Corollary 4.1. *Given a seller* i *and a buyer* k*, there exists* δ^**, such that for all* $\delta > \delta^*$ *and at any equilibrium the following is true. If* $\lambda_{ij} > 0$ *and* $\lambda_{jk} > 0$ *for a middlemen* j*, that is trade occurs along the route* $i \to j \to k$*, then the cost* $C_{ij} + C_{jk}$ *is the smallest among all trading routes between* i *and* k*.*

[4] Again we drop the explicit dependence on δ for the payoffs, trade probabilities and stationary state of the middlemen.

Remark: The above result demonstrates a global-level efficiency that emerges in the equilibria of the local non-cooperative bargaining scheme if agents are patient enough: edges that are not along a cheapest path from any seller and buyer pair are never used, and middlemen who have no edges along a cheapest path from any seller and buyer pair see no trade. Note that there is no claim that either all or even any of the cheapest path routes between a given seller and buyer pair are used in an equilibrium.

Proof. Consider equations (10), (11), (12) and (13). As δ approaches 1, the $\log(1/\delta)$ term approaches 0. Since $u_1(j), u_0(j) \in [0, \max_{k\in\mathcal{K}} V_k]$ for all $j \in \mathcal{J}$, $u_1(i) \in [0, \max_{k\in\mathcal{K}} V_k]$ for all $i \in \mathcal{I}$, and $u_0(k) \in [0, \max_{k\in\mathcal{K}} V_k]$ for all $k \in \mathcal{K}$, it has to be that given any $\epsilon > 0$, there exists δ^* such that for all $\delta > \delta^*$, we have

$$\begin{aligned} z_{ik} &\le \epsilon \quad \forall (i,k) \in \mathcal{E}_1, \\ z_{ij} &\le \epsilon \quad \forall (i,j) \in \mathcal{E}_2, \\ z_{jk} &\le \epsilon \quad \forall (j,k) \in \mathcal{E}_3. \end{aligned}$$

Now consider a pair of agents, one seller i and buyer k. We have three cases then:

(1) All trade routes from i to k have to visit some middleman. Let $j \in \mathcal{J}$ be one such middleman so that $(i,j) \in \mathcal{E}_2$ and $(j,k) \in \mathcal{E}_3$. The inequalities above then imply the following:

$$\begin{aligned} u_1(j) - u_0(j) &\ge u_1(k) - u_0(k) - C_{jk} - \epsilon, \\ u_1(j) - u_0(j) &\le u_1(i) - u_0(i) + C_{ij} + \epsilon. \end{aligned}$$

These with $u_0(i) = 0$ and $u_1(k) = V_k$ imply

$$u_1(i) + u_0(k) \ge V_k - C_{ij} - C_{jk} - 2\epsilon.$$

Note that this inequality holds for every $j \in \mathcal{J}$ that lies along a trade route from i to k. Therefore,

$$u_1(i) + u_0(k) \ge V_k - \min_{\{j:(i,j)\in\mathcal{E}_2 \text{ and } (j,k)\in\mathcal{E}_3\}} \Big(C_{ij} + C_{jk} \Big) - 2\epsilon.$$

Because ϵ can be chosen arbitrarily small, thus for any middleman j who is not on a smallest transaction cost path from i to k, we can choose δ close enough to 1 such that either z_{ij} or z_{jk} is strictly negative and so no trade can occur on the corresponding edge;

(2) Notice that the same argument also works for the case, where if in addition to the middlemen, there also exists a direct link between i and k. Then

$$u_1(i) + u_0(k) \ge V_k - \min\left(C_{ik}, \min_{\{j:(i,j)\in\mathcal{E}_2 \text{ and } (j,k)\in\mathcal{E}_3\}} \Big(C_{ij} + C_{jk} \Big) \right).$$

Again it is clear that no trade occurs over links that are not part of a smallest transaction cost path from i to k;

(3) If i and k only have a direct route between them, then that is the only route via which trade can occur between this seller and buyer pair. Also, if no routes exist between i and k, then obviously no trade occurs between these two agents.

□

A refined analysis is needed to distinguish between multiple cheapest routes in order to determine which ones get used; the division of the transaction cost, in particular, the values closer to the seller, allow us to differentiate between the multiple cheapest routes.

4.2. **Endogenous Delay.** We now consider a simple network that consists of two links. This network represents the simplest example where sellers and buyers cannot trade directly. We fully characterize the limit stationary equilibrium in this example, which will be shown to be unique. Even in this simple network, we observe an interesting phenomenon of endogenous delay as part of the equilibrium. This is counterintuitive since in a full information dynamic bargaining model like ours, delay in trade does not enable agents to learn any new information, but only decreases the total surplus of trade. Therefore, the network structure and the combination of incentives of long-lived and short-lived agents are the main sources causing this inefficiency in bargaining.

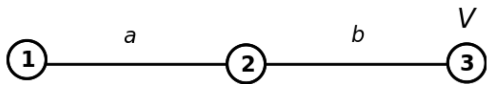

FIGURE 2. Network

Assume a, b are transaction costs of the first and second link, also let V be the value of the consumption of the good; without loss of generality we will insist that trade is favorable so that $V > a + b$. Abusing notation we will index the links as a and b. The probability of using the links is then (π_a, π_b). We assume the population sizes at every node is equal, and without loss of generality, we assume $N_1 = N_2 = N_3 = 1$. From this point onwards, we will also assume that we start with exactly one agent at each location [5]. We will show that in this simple network, the stationary equilibrium is unique, and we characterize the condition on which agents do not trade immediately.

Theorem 4.2. *When the agents are patient, there is always a unique limit stationary equilibrium. Furthermore, if* $V \geq \left(1 + \frac{\pi_a}{\pi_a + \pi_b}\right) a + b =: \bar{V}$*, then trade always happen, otherwise there is a delay. The probability of trade, the payoffs and the equilibrium state of the middleman are given by*

$$\lambda_b = 1, \quad \mu_j = \begin{cases} \frac{\pi_a}{\pi_a + \pi_b} & \textit{if } V \geq \bar{V} \\ \frac{V-b-a}{a} & \textit{otherwise} \end{cases}, \qquad \lambda_a = \begin{cases} 1 & \textit{if } V \geq \bar{V} \\ \frac{\pi_b(V-b-a)}{\pi_a(2a+b-V)} & \textit{otherwise} \end{cases},$$

$$u_1(j) = \begin{cases} \frac{(V-b)\left(2 - \frac{\pi_a}{\pi_a+\pi_b}\right) - a}{1 + \frac{\pi_a \pi_b}{(\pi_a+\pi_b)^2}} & \textit{if } V \geq \bar{V} \\ a & \textit{otherwise} \end{cases}, \qquad u_0(j) = \begin{cases} \frac{V - \left(1 + \frac{\pi_a}{\pi_a+\pi_b}\right)a - b}{1 + \frac{\pi_a \pi_b}{(\pi_a+\pi_b)^2}} & \textit{if } V \geq \bar{V} \\ 0 & \textit{otherwise} \end{cases},$$

$$u_0(k) = \begin{cases} \frac{(V-b)\left(2 - \frac{\pi_a}{\pi_a+\pi_b}\right) - a}{\frac{\pi_a+\pi_b}{\pi_a} + \frac{\pi_b}{\pi_a+\pi_b}} & \textit{if } V \geq \bar{V} \\ 1 - b - a & \textit{otherwise} \end{cases}, \qquad u_1(i) = \begin{cases} \frac{V - \left(1 + \frac{\pi_a}{\pi_a+\pi_b}\right)a - b}{\frac{\pi_a+\pi_b}{\pi_a} + \frac{\pi_b}{\pi_a+\pi_b}} & \textit{if } V \geq \bar{V} \\ 0 & \textit{otherwise.} \end{cases}$$

Remark: Note that trade always occurs on link b but can be delayed at link a. Since the buyer is at the other end of link b, it stands to reason that there is no delay in the trade. However, at link a, any sale of the item results in a decreased likelihood of the trade at the same link (in the near future) and this opportunity cost introduces the delay in trade. Note also that with a delay in trade, the seller obtains no surplus! We will revisit this effect in the next section. Trade gets delayed when the value of the good is below a specific threshold. From the proof one can discern that the additional penalty term in the threshold is the product of the transaction cost at link a and the stationary probability that the middleman possesses the good. This can be viewed as a carryover of the sunk cost at the second stage and should be contrasted with the result in [14] where the additional penalty without search friction is the entire transaction cost at link a.

Proof. See Appendix A.3 □

[5]Following the proofs of our results, it will become clear that this assumption does not result in any loss of generality when agents are patient.

4.3. **Fairness.** Lastly, we consider the imbalance between the surplus of sellers and buyers as a result of our decentralized trade model.

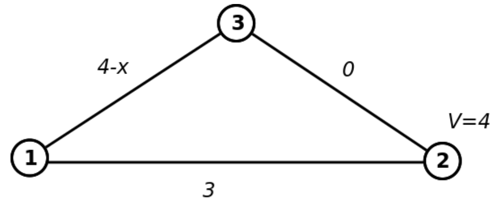

FIGURE 3. Network

Consider the following (also simple) network, where node 1 represents sellers, node 2 represent buyers and node 3 represents middlemen, as illustrated in Figure 3. Again without loss of generality, we assume that$N_1 = N_2 = N_3 = 1$. We also assume that in our bargaining model, every link is selected uniformly at random, that is $\pi_{ij} = 1/3$ for all $i \neq j$. Assume the buyer's valuation for the good is $V_2 = 4$, and the transaction costs are the following: $C_{12} = 3$, $C_{32} = 0$ and $C_{13} = 4 - x$. We will investigate the equilibrium as x changes. As x increases, the transaction cost between 1 and 3 decreases, making the total trade surplus $\max\{4-3, 4-(4-x)\} = \max\{1, x\}$ increase.

The surplus of sellers in this example is understood as the payoff of agent at node 1 when owning an item: $u_1(1)$. On the other hand, the surplus of buyers in this example is the payoff of agent at node 2 when not owning an item: $u_0(2)$. According to the analysis in Section 4.1, as the discount rate δ approaches 1 trade will only goes through the cheapest route. Let $\bar{C}$ be the cost of this route, as seen in Section 4.1 we also have

$$lim_{\delta\to 1} u_1(2) - u_0(2) - (u_1(1) - u_0(1)) = \bar{C}.$$

This is equivalent to

$$lim_{\delta\to 1} u_0(2) + u_1(1) = V_2 - \bar{C}.$$

In other words, the total surplus of a seller and a buyer approaches the total trading surplus. This seems counter-intuitive at first, because this means for every transaction, middlemen only make a vanishing amount of fee. However, this effect is due to the fact that sellers and buyers are short-lived, while middlemen are long-lived, and thus middlemen can earn a positive payoff by accumulating fees over an infinite horizon.

Now, in the example above, when considering the equilibrium payoff as δ approaches 1, we have if $x < 1$, sellers and buyers will trade directly, and in this case sellers and buyers equally share the surplus, in which case their surplus is $\frac{4-3}{2} = \frac{1}{2}$.

On the other hand, if $x > 1$, then direct trade between sellers and buyers is too expensive, trade will go through middlemen at node 3. In the latter case, we will use the analysis in Section 4.2 to compute the equilibrium payoff, and we have

(1) $1 < x < 4/3$: Seller's surplus, $u_1(1) = 0$ and buyer's surplus $u_0(2) = x$, so that the buyers get all the trade surplus;
(2) $4/3 \leq x \leq 4$: Seller's surplus, $u_1(1) = \frac{4-3/2(4-x)}{5/2} = \frac{3x-4}{5}$ and buyer's surplus $u_0(2) = \frac{2x+4}{5}$.

(See Figure 4).

Even in this simple network, we observe quite an interesting phenomenon on the discontinuous shift in the trading pattern occurring in the network. If the transaction cost $C_{13} = 4 - x$ between 1 and 3 decreases the total surplus between sellers and buyers increases, but sellers are actually worse off because of this shift in the market structure. This also highlights how local adjustments by the sellers could leave them in a worse-off position.

This example captures an interesting and counterintuitive phenomenon: as the transaction cost towards middlemen decreases, sellers can be worse off because of high cost in direct trading, buyers

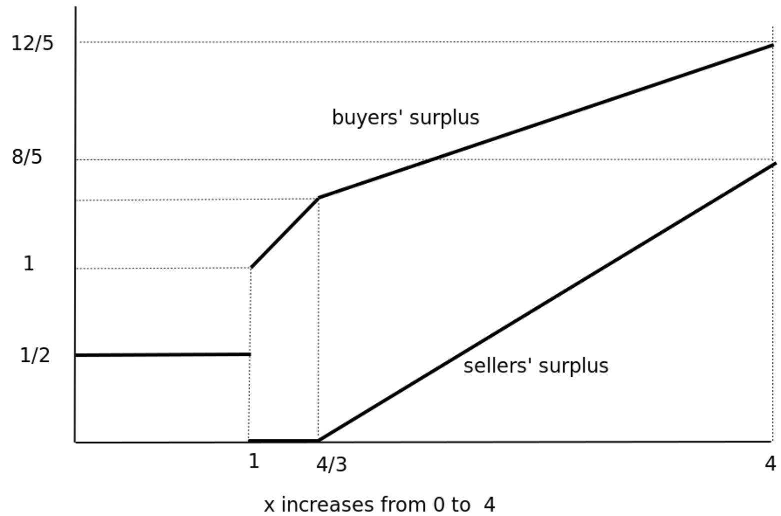

FIGURE 4. Surplus of seller and buyers as x increases from 0 to 4

refuse to trade directly and prefer to trade through middlemen. For example, in many supply chain networks, as these global networks get large, sellers and buyers do not trade directly and several types of organizations emerge as middlemen. In many cases such as in coffee industry, sellers (coffee farmers) obtain a very small fraction of surplus because there are too many middlemen in the supply chain network. See for example [2] for a related empirical analysis of the coffee global supply chain and the recent shift in its market structure.

5. Conclusions and Future Work

In this paper we considered non-cooperative local bargaining over a trading network with a single type of good. In the limiting scenario of many agents, we showed the existence of a limit stationary equilibrium that can be characterized by a combination of the stationary probability of a trade happening on each link and the stationary distribution of the good at the agents. We then showed that when agents are patient enough, this limiting equilibrium can exhibit global efficiency. We applied this concept to several simple network structures to study the impact of the network on the bargaining power and surplus of all agents. In future work we plan to extend the results to more general networks, to the trade of multiple goods, and to the analysis of risk of losing or damaging the good *en route*.

References

[1] Adlakha, S., Johari, R., and Weintraub, G. Y. 2010. Equilibria of dynamic games with many players: Existence, approximation, and market structure. http://arxiv.org/abs/1011.5537.

[2] Bacon, C. 2005. Confronting the coffee crisis: Can fair trade, organic, and specialty coffees reduce small-scale farmer vulnerability in northern Nicaragua? *World Development* 33.

[3] Benaim, M. and Boudec, J.-Y. L. 2011. On mean field convergence and stationary regime. http://arxiv.org/abs/1111.5710.

[4] Blume, L., Easley, D., Kleinberg, J., and Tardos, E. 2009. Trading networks with price-setting agents. *Games and Economic Behavior 67,* 1, 36–50.

[5] Ethier, S. N. and Kurtz, T. G. 2005. *Markov processes: Characterization and convergence* 2nd Ed. John Wiley & Sons.

[6] Graham, C. and Méléard, S. 1994. Chaos hypothesis for a system interacting through shared resources. *Probability Theory and Related Fields 100,* 2, 157–174.

[7] Graham, C. and Méléard, S. 1997. Stochastic particle approximations for generalized Boltzmann models and convergence estimates. *The Annals of probability 25,* 1, 115–132.

[8] Guéant, O., Lasry, J., and Lions, P. 2011. Mean field games and applications. *Paris-Princeton Lectures on Mathematical Finance 2010*, 205–266.

[9] Kleinberg, J. and Tardos, É. 2008. Balanced outcomes in social exchange networks. In *Proceedings of the 40th annual ACM symposium on Theory of computing.* ACM, 295–304.

[10] Kreps, D. 1977. A note on 'fulfilled expectations' equilibria. *Journal of Economic Theory 14*, 32–43.

[11] Lasry, J. and Lions, P. 2007. Mean field games. *Japanese Journal of Mathematics 2,* 1, 229–260.

[12] Manea, M. 2011. Bargaining in stationary networks. *The American Economic Review 101,* 5, 2042–2080.

[13] Manea, M. 2012. Bargaining in dynamic markets with multiple populations. Tech. rep., Mimeo.

[14] Nguyen, T. 2012. Local bargaining and endogenous fluctuations. In *Proc. 13th ACM Conf. on Electronic Commerce (EC-2012).* Valencia, Spain.

[15] Rubinstein, A. and Wolinsky, A. 1987. Middlemen. *The Quarterly Journal of Economics 102,* 3, 581–593.

[16] Wong, Y.-Y. and Wright, R. 2011. Buyers, sellers and middlemen: Variations on search-theoretic themes. Working Paper. Working Paper.

Appendix A. APPENDIX

A.1. **Proof of Theorem 3.2.** By an application of Kurtz's Theorem [5, Th. 2.1, Chapter 11], we obtain a differential equation for the limiting process that is a continuous function from the non-negative reals to $\prod_{i\in\mathcal{I}}[0,N_i]\times\prod_{j\in\mathcal{J}}[0,N_j]\times\prod_{k\in\mathcal{K}}[0,N_k]$. The limiting processes[6] are given as follows for all $t\geq 0$,

$$\begin{aligned}
&\forall i\in\mathcal{I},\quad x_i(t)\equiv N_i;\qquad\qquad \forall k\in\mathcal{K},\quad x_k(t)\equiv 0;\\
&\forall j\in\mathcal{J},\quad \frac{dx_j(t)}{dt}=\left(1-\frac{x_j(t)}{N_j}\right)\sum_{i\in\mathcal{I}:(i,j)\in\mathcal{E}_2}\pi_{ij}\lambda_{ij}-\frac{x_j(t)}{N_j}\sum_{k\in\mathcal{K}:(j,k)\in\mathcal{E}_3}\pi_{jk}\lambda_{jk}\\
&\qquad\qquad\qquad\qquad=\sum_{i\in\mathcal{I}:(i,j)\in\mathcal{E}_2}\pi_{ij}\lambda_{ij}-\frac{x_j(t)}{N_j}\left(\sum_{i\in\mathcal{I}:(i,j)\in\mathcal{E}_2}\pi_{ij}\lambda_{ij}+\sum_{k\in\mathcal{K}:(j,k)\in\mathcal{E}_3}\pi_{jk}\lambda_{jk}\right).
\end{aligned}$$

Using a quadratic Lyapunov function (square of the distance to the equilibrium point) it follows that there is a unique and globally asymptotically stable equilibrium point that is given by

$$\begin{aligned}
&\forall i\in\mathcal{I},\quad x_i^*=N_i;\qquad\qquad \forall k\in\mathcal{K},\quad x_k^*=0;\\
&\forall j\in\mathcal{J},\quad x_j^*=N_j\frac{\sum_{i\in\mathcal{I}:(i,j)\in\mathcal{E}_2}\pi_{ij}\lambda_{ij}}{\sum_{i\in\mathcal{I}:(i,j)\in\mathcal{E}_2}\pi_{ij}\lambda_{ij}+\sum_{k\in\mathcal{K}:(j,k)\in\mathcal{E}_3}\pi_{jk}\lambda_{jk}}.
\end{aligned}$$

Therefore, the fraction of agents with the good satisfies

$$\begin{aligned}
&\forall i\in\mathcal{I},\quad \mu_i=1;\qquad\qquad \forall k\in\mathcal{K},\quad \mu_k=0;\\
&\forall j\in\mathcal{J},\quad \mu_j=j\frac{\sum_{i\in\mathcal{I}:(i,j)\in\mathcal{E}_2}\pi_{ij}\lambda_{ij}}{\sum_{i\in\mathcal{I}:(i,j)\in\mathcal{E}_2}\pi_{ij}\lambda_{ij}+\sum_{k\in\mathcal{K}:(j,k)\in\mathcal{E}_3}\pi_{jk}\lambda_{jk}}.
\end{aligned}$$

[6]Even though Kurtz's Theorem applies only for finite time horizons, the compact setting of the scaled processes and the well-behaved nature of the differential equation above allow us to analyze the converges of the stationary solutions as well.

For each m, it is easy to see that the Markov process is irreducible and has finite states, and so is positive recurrent. Thus, owing to the compact setting, the stationary measures of the scaled state processes converge to the point mass on the equilibrium point as m increases without bound, see [3].

A.2. **Proof of Theorem 3.3.** We need to show that there exists $(\lambda, \mu, \vec{u})$ satisfying the following conditions

(1) Convergence: given the trading dynamics defined by λ, the replicated economy converges to the steady state μ. According to Theorem 3.2, this is equivalent to the condition (17)
(2) Payoff-state consistency: $\lambda, \mu, \vec{u}$ need to satisfy the ..
(3) Payoff-dynamic consistency: if $z_{ij} > 0$ then $\lambda_{ij} = 1$; if $z_{ij} < 0$ then $\lambda_{ij} = 0$; and if $0 < \lambda_{ij} < 1$, then $z_{ij} = 0$

$$F(\lambda, \mu, u) = (\Lambda, \mu', u'),$$

where

$$\begin{aligned}
\forall i \in \mathcal{I} \quad & \mu'_i = 1, \\
\forall k \in \mathcal{K} \quad & \mu'_k = 0, \\
\forall j \in \mathcal{J} \quad & \mu'_j = \frac{\sum_{i \in \mathcal{I}:(i,j) \in \mathcal{E}_2} \pi_{ij} \lambda_{ij}}{\sum_{i \in \mathcal{I}:(i,j) \in \mathcal{E}_2} \pi_{ij} \lambda_{ij} + \sum_{k \in \mathcal{K}:(j,k) \in \mathcal{E}_3} \pi_{jk} \lambda_{jk}},
\end{aligned}$$

and

$$\begin{aligned}
\Lambda_{ij} =& \{1\} \text{ if } u_1(j) - u_0(j) - (u_1(i) - u_0(i)) - C_{ij} > 0, \\
\Lambda_{ij} =& \{0\} \text{ if } u_1(j) - u_0(j) - (u_1(i) - u_0(i)) - C_{ij} < 0, \\
\Lambda_{ij} =& [0, 1] \text{ if } u_1(j) - u_0(j) - (u_1(i) - u_0(i)) - C_{ij} = 0.
\end{aligned}$$

Furthermore,

$$\begin{aligned}
u'_0(i) =& 0 \text{ for all sellers } i \in \mathcal{I}, \\
u'_1(i) =& \sum_{j \in \mathcal{J} \cup \mathcal{K}} \frac{\pi_{ij}}{2N_i \ln(1/\delta))} (1 - \mu_j) \max\{u_1(j) - u_0(j) - (u_1(i) - u_0(i)) - C_{ij}, 0\} \text{ for all sellers } i \in \mathcal{I}, \\
u'_1(k) =& V_k \text{ for all buyers } k \in \mathcal{K}, \\
u'_0(k) =& \sum_{j \in \mathcal{J} \cup \mathcal{I}} \frac{\pi_{jk}}{2N_k \ln(1/\delta))} \mu_j \max\{u_1(k) - u_0(k) - (u_1(j) - u_0(j)) - C_{jk}, 0\} \text{ for all buyer } k \in \mathcal{K}, \\
u'_0(j) =& \sum_{i \in \mathcal{I}} \frac{\pi_{ij}}{2N_j \ln(1/\delta))} \mu_i \max\{u_1(j) - u_0(j) - (u_1(i) - u_0(i)) - C_{ij}, 0\} \text{ for all middlemen } j \in \mathcal{J}, \\
u'_1(j) =& \sum_{k \in \mathcal{K}} \frac{\pi_{jk}}{2N_j \ln(1/\delta))} (1 - \mu_k) \max\{u_1(k) - u_0(k) - (u_1(j) - u_0(j)) - C_{jk}, 0\} \text{ for all middlemen } j \in \mathcal{J}.
\end{aligned}$$

It is straightforward to check that the function above satisfies all the requirements for Kakutani's fixed-point theorem: the domain is a non-empty, compact and convex subset, the mapping is an upper hemicontinuous set-valued function and the image of any point in the domain is non-empty, closed and convex. Therefore, there must be a fixed-point, and furthermore, by definition, any fixed point of this mapping is a limit stationary equilibrium.

A.3. **Proof of Theorem 4.2.** The equilibrium equations for this case are as follows:

$$u_1(i) = \frac{\pi_a}{2\ln(1/\delta)}(1-\mu_j)\max\{z_a, 0\}, \qquad u_0(j) = \frac{\pi_a}{2\ln(1/\delta)}\max\{z_a, 0\},$$
$$u_1(j) = \frac{\pi_b}{2\ln(1/\delta)}\max\{z_b, 0\}, \qquad u_0(k) = \frac{\pi_b}{2\ln(1/\delta)}\mu_j\max\{z_b, 0\},$$
$$z_a = \delta(u_1(j) - u_0(j) - u_1(i)) - a, \qquad z_b = \delta(V - u_0(k) - u_1(j) + u_0(j)) - b,$$
$$\lambda_a \in \begin{cases} \{1\} & z_a > 0 \\ \{0\} & z_a < 0 \\ [0,1] & z_a = 0 \end{cases}, \quad \lambda_b \in \begin{cases} \{1\} & z_b > 0 \\ \{0\} & z_b < 0 \\ [0,1] & z_b = 0 \end{cases}, \qquad \mu_j = \frac{\pi_a\lambda_a}{\pi_a\lambda_a + \pi_b\lambda_b}.$$

From the above is clear that $u_1(i) = (1-\mu_j)u_0(j)$ and $u_0(k) = \mu_j u_1(j)$. Substituting these we get

$$z_a = \delta\big(u_1(j) - (2-\mu_j)u_0(j)\big) - a, \qquad z_b = \delta\big(V - (1+\mu_j)u_1(j) + u_0(j)\big) - b.$$

First consider the assumption that trade occurs with probability one on both links, i.e., $\lambda_a = \lambda_b = 1$. This then implies that $\mu_j = \frac{\pi_a}{\pi_a+\pi_b}$, $z_a, z_b \geq 0$ and we can substitute them directly into the equations for the payoffs. We then obtain the following linear equations in $u_0(j)$ and $u_1(j)$,

$$u_0(j) = \frac{\delta u_1(j) - a}{\frac{2\ln(1/\delta)}{\pi_a} + \delta(2-\mu_j)}, \quad u_1(j) = \frac{\delta u_0(j) + \delta V - b}{\frac{2\ln(1/\delta)}{\pi_b} + \delta(1+\mu_j)}.$$

We can take limits in the equations above as δ goes to 1 (along an appropriate subsequence) to get

$$u_0(j)(2-\mu_j) = u_1(j) - a, \qquad u_1(j)(1+\mu_j) = u_0(j) + V - b.$$

The unique solution is

$$u_1(j) = \frac{(2-\mu_j)(V-b) - a}{1+\mu_j-\mu_j^2}, \quad u_0(j) = \frac{V - (1+\mu_j)a - b}{1+\mu_j-\mu_j^2}. \tag{20}$$

It is easily seen that $u_1(j) \geq 0$ and $u_0(j) \geq 0$ if and only if $V \geq (1+\mu_j)a - b = \bar{V}$, and at the equilibrium $z_a = z_b = 0$.

For the remainder assume that $V < \bar{V}$. Consider the case that $\lambda_b = 1$ and $0 < \lambda_a < 1$. This then implies that $z_a = 0$, $z_b \geq 0$ and $\mu_j = \frac{\pi_a\lambda_a}{\pi_a\lambda_a+\pi_b}$. Again taking a limit of δ going to 1 (along an appropriate subsequence), we also get $z_b = 0$. If we solve (20), then the calculated $u_0(j)$ will be negative which then implies that at the equilibrium $u_0(j) = 0$; note that $z_a = 0$ for $\delta < 1$ also yields the same conclusion. Now it follows that

$$u_1(j) = a, \; \mu_j = \frac{V-b-a}{a}, \text{ and } \lambda_a = \frac{\pi_b(V-b-a)}{\pi_a(2a+b-V)} \in (0,1).$$

Since $V < \bar{V} = \left(1 + \frac{\pi_a}{\pi_a+\pi_b}\right)a + b$, it also follows that $V < 2a + b$ which ensures $\mu_j \leq 1$. Similarly, one can verify that $\lambda_a \in (0,1)$. Since the consistency conditions are met, we have an equilibrium.

The uniqueness of the solution in both cases also proves that the same solution holds along every subsequence of δ converging to 1 (from below) so that the uniqueness of the equilibrium also follows. Finally, we can verify that there can be no other equilibria.